

EP1033650

Publication Title:

Instruction translation method

Abstract:

A method of translating source code instructions into target code instructions is described. Prior to translate time, an existing interpreter is analysed to identify sequences that implement individual source order code instructions. Subsequences within each template that implement predetermined sub-functions are identified and eliminated. The sequences are compiled and stored as templates. For each instruction in an input block of source code instructions, the appropriate template for that source code instruction is selected and appended to an output block of target code instructions. The source code block is then analysed to determine the net effect of the non-implemented sub-functions, and code is planted in the output block to achieve this net effect.

Data supplied from the esp@cenet database - <http://ep.espacenet.com>

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 033 650 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
06.09.2000 Bulletin 2000/36

(51) Int Cl.7: G06F 9/45, G06F 9/455

(21) Application number: 00300485.0

(22) Date of filing: 24.01.2000

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Hughes, Kevin
Glossop, Derbyshire SK13 8RJ (GB)
• Eaton, John Richard
Salford M6 8EY (GB)

(30) Priority: 04.03.1999 GB 9904823

(74) Representative: Guyatt, Derek Charles et al
International Computers Limited
Intellectual Property Department
Cavendish Road
Stevenage, Herts, SG1 2DY (GB)

(71) Applicant: International Computers Ltd.
London, EC2A 1SL (GB)

(54) Instruction translation method

(57) A method of translating source code instructions into target code instructions is described. Prior to translate time, an existing interpreter is analysed to identify sequences that implement individual source order code instructions. Subsequences within each template that implement predetermined sub-functions are identified and eliminated. The sequences are compiled

and stored as templates. For each instruction in an input block of source code instructions, the appropriate template for that source code instruction is selected and appended to an output block of target code instructions. The source code block is then analysed to determine the net effect of the non-implemented sub-functions, and code is planted in the output block to achieve this net effect.

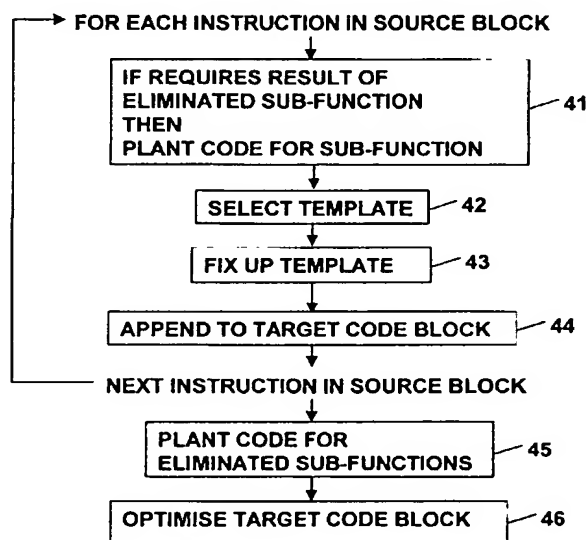


FIG.4

EP 1 033 650 A2

Description

Background to the Invention

[0001] This invention relates to a method for translating instructions in a computer system.

[0002] The invention is particularly concerned with a computer system in which source code instructions are translated into target code instructions for execution on a particular processor. This may be required, for example, where one processor is being used to emulate another, in which case the instructions for the processor being emulated must be translated into instructions for the emulating processor.

[0003] One approach, referred to as interpretation, is to create a software model of the processor being emulated. This model operates by reading each target instruction, decoding it, and selecting one of a number of sequences that perform the same function as the instruction being emulated. This fetch/decode/execute sequence is repeated for each source code instruction in turn.

[0004] A more efficient approach is to translate a block of source code instructions, rather than a single instruction. That is, the source code is divided into blocks, and each source code block is translated into a block of target code instructions, functionally equivalent to the source code block. Typically, a block has a single entry point and one or more exit points. The entry point is the target of a source code jump, while the (or each) exit is a source code jump.

[0005] Translating blocks is potentially much more efficient, since it provides opportunities for eliminating redundant instructions within the target code block, and other optimisations. Known optimising compiler techniques may be employed for this purpose. To increase efficiency further, the target code blocks may be held main memory and/or a cache store, so that they are available for re-use if the same section of code is executed again, without the need to translate the block.

[0006] However, the process of designing such a translator is very complex, and it is difficult to avoid errors in translation. One object of the present invention is to provide an improved translation technique, which reduces such errors.

Summary of the Invention

[0007] According to the invention, in a computer system, a method of translating source code instructions into target code instructions comprises the steps:

(a) analysing an existing interpreter to identify sequences that implement individual source code instructions and storing those sequences as templates; and

(b) for each instruction in an input block of source

code instructions, selecting an appropriate template for that source code instruction and appending this template to an output block of target code instructions.

[0008] It can be seen that this uses an existing interpreter as the basis for building a translation mechanism. Assuming that the existing interpreter is already fully validated, the possibility of errors in the templates is correspondingly reduced. In effect, the invention provides a way of "leveraging" an existing interpreter.

[0009] One embodiment of the invention will now be described by way of example with reference to the accompanying drawings.

Brief Description of the Drawings

[0010] Figure 1 is a block diagram showing the main data structures and processes involved in this embodiment.

[0011] Figure 2 is a flow chart of a process for forming templates from an existing interpreter code.

[0012] Figure 3 is a flow chart of a process for initialising the templates.

[0013] Figure 4 is a flow chart of a process for translating a block of source code, using the templates.

Description of an Embodiment of the Invention

[0014] The present embodiment is concerned with a mechanism for translating instructions from a source instruction set into a target instruction set. For example, the source instruction set may be the ICL VME instruction set, and the target instruction set may be a micro-processor assembler language.

[0015] The ICL VME instruction set is complex, and each instruction may involve a number of sub-functions such as, for example:

- calculating an operand address,
- performing range checks on the operand address; for example, checking that the address of a stack operand is not greater than the current stack front,
- fetching or storing an operand,
- writing or reading an operand to or from a register,
- clearing or setting an overflow register OV,
- incrementing a program counter register PC by an amount dependent on the instruction length.

[0016] It is assumed that there exists a fully validated interpreter for translating all instructions in the source instruction set.

[0017] This interpreter may be written either in a high-level language, as a set of macros, or in assembler language, with a defined sequence for each instruction in the source instruction set. Each of these interpreter sequences, in turn, includes a number of sub-sequences, for translating the sub-functions in the instruction. For

example, an interpreter sequence may contain sub-sequences for calculating the operand address, performing range checks on the operand address, and so on.

[0018] Referring to Figure 1, prior to system build time, a Create Templates process 10 is performed. This process takes the existing interpreter source code 11, and generates a set of template source code sequences 12, one for each instruction in the source instruction set. The template source code sequences at this stage are source code sequences derived from the interpreter source code. In the present embodiment of the invention, the process 10 is performed manually, but in other embodiments it could be performed automatically, by a suitable program.

[0019] Figure 2 is a flowchart of the Create Templates process 10. In this process, the source interpreter code 11 is scanned, to identify the sequences corresponding to individual instructions in the source instruction set. For each sequence, the following actions are performed.

[0020] (Step 21) First, the interpreter sequence is copied to the template source code 12. The template source code is then scanned to detect sub-sequences representing certain predetermined common sub-functions, and these sub-functions are removed from the template code. In this example, the sub-functions that are removed include PC update, OV clearing, and operand address range checks.

[0021] In the case where the interpreter code is macro-generated, the removal of the predetermined sub-functions is achieved by modifying the macros. For example, consider the case where the interpreter contains the following macros:

- CHECK_VA(x,y) - check a virtual address x against a limit y.
- CLEAROV() - clear the overflow register OV.
- UPDATEPC(n) - add n to the program counter register PC. In this case, CHECK_VA(x,y) would be modified to always return "true" (if this check can be performed statically), while CLEAROV() and UPDATEPC(n) would be modified to do nothing.

[0022] (Step 22) Next, the template source code 12 is scanned, looking for constants that will be derived from literal values in the target instructions at translate (run) time. These constants are replaced with predetermined literal marker values.

[0023] (Step 23) Finally, the start and end of each template source code sequence are marked by planting suitable binary codes in the template source code 12.

[0024] Referring again to Figure 1, the resulting template source code sequences 12 are then compiled as part of the standard system build process 13, resulting in a set of binary templates 14.

[0025] At system initialisation, an Initialise Templates process 15 is performed. Referring to Figure 3, this process performs the following actions for each binary tem-

plate 14 in turn.

[0026] (Step 31) First, the Initialise Templates process scans the binary template to locate the start/end of the binary template (using the binary codes planted at step 23 above). This information is added to a set of data structures 16 to enable the translator to locate the binary template.

[0027] (Step 32) Next, the binary template is scanned to locate the marker values that were inserted at step 22 above, and to locate all calls within the template.

[0028] (Step 33) For each marker value in the template, the Initialise Templates process inserts an entry (referred to herein as a "fix-up" entry) in the data structures 16. The fix-up entry identifies the location of the marker value and specifies the data type of the constant value that is to be inserted into the code at translate time (run time) to replace the marker value. Similarly, for each call in the template, the Initialise Templates process inserts a fix-up entry in the data structures 16, identifying the location of the call.

[0029] At run time, a Translation process 17 (Figure 1) is performed. This process takes source code blocks 18, and translates each block into a target code block 19, using the binary templates 14 and the data structures 16. Referring to Figure 4, the process 17 selects each instruction in the source code block 18 in turn, and performs the following actions on the currently selected instruction.

[0030] (Step 41) First, the Translation process determines whether the current instruction requires an operand that would have been generated by one of the eliminated sub-functions. If so, it plants code in the target code block 19, to perform the actions of the eliminated sub-sequence. For example, suppose the Translator procedure finds that the current instruction uses the value of the PC register, but one or more preceding instructions in the block have had their PC update sub-functions eliminated. In this case, the Translator procedure will plant code to bring the PC register up to date.

[0031] (Step 42) Next, the Translation process identifies the instruction type of the current instruction, and selects the binary template 14 corresponding to this instruction type from the data structure created in step 31.

[0032] (Step 43) The Translation process then looks up the data structures 16, to identify any fix-up entries for the template.

[0033] In the case where a fix-up entry represents a constant, the process inserts the required constant. The constant is derived from the parameters of the current instruction, according to its specified data type. In the case where a fix-up entry represents a call, the process inserts the required information for the call.

[0034] (Step 44) The fixed-up template is appended to the output target code block 19.

[0035] (Step 45) After all the instructions in the source code block have been processed in this way, the Translation process scans the source code block 18 to determine the net effect of the eliminated sub-functions. It

then plants code in the target code block 19, to ensure that this net effect is achieved.

[0036] In this step, the Translation process determines which of the instructions in the source code block require an operand address range check. It then plants code at the beginning of the target code block 19, to perform a combined range check, having the same net effect as all the eliminated range checks in the block.

[0037] For example, consider the following block of source code:

```
N=5    check 5 < SF-LNB
N=6    check 6 < SF-LNB
N=7    check 7 < SF-LNB
```

(where SF denotes a stack front register and LNB denotes a local name base register). The sub-functions for these three range checks will be eliminated by the Create Templates process. The Translation process therefore plants a single merged range check, which checks SF-LNB against the maximum of the three values (5,6,7) :

7 < SF-LNB. It can be seen that if this test passes, all three original tests must pass.

[0038] Now consider the following block of code:

```
N=5    check 5 < SF(0)-LNB
SF(1)=SF(0)+5
N=11   check 11 < SF(1)-LNB
SF(2)=SF(1)-5
N=7    check 7 < SF(2)-LNB
```

In this case the value of SF is modified during execution of the block, and hence a merged check based on the maximum of the three values (5,11,7) would fail.

[0039] The Translation process solves this problem by tracking the adjustment of SF, and replacing the checks as follows:

```
N=5    check 5 < SF(0)-LNB
SF(1)=SF(0)+5
N=11   check 11 < SF(0)+5-LNB
SF(2)=SF(1)-5
N=7    check 7 < SF(0)+5-5-LNB
```

[0040] This is simplified to become:

```
N=5    check 5 < SF(0)-LNB
SF(1)=SF(0)+5
N=11   check 6 < SF(0)-LNB
SF(2)=SF(1)-5
N=7    check 7 < SF(0)-LNB
```

The translation process then plants a single check, based on the maximum of these new values i.e. (7 < SF-LNB). This approach works both for positive and negative adjustment of SF.

[0041] Also in step 45, the Translation process deter-

mines the amount by which each instruction in the source code block updates the PC register, and then plants a single instruction at the end of the target code block, to increment the PC register by the total of all these updates.

[0042] (Step 46) Finally, as an optional step, the Translation process may perform further optimisations on the target code block 19, using conventional optimisation techniques, such as register tracking to eliminate redundant register reads and writes.

[0043] In summary, it can be seen that the translation mechanism described above uses an existing interpreter to form templates for translating individual instructions. Because the interpreter is fully validated, the templates should also be error-free.

[0044] The efficiency of the generated code is improved by eliminating certain common sub-functions (such as "update PC") from the templates, and planting code in the target code block to restore the net effect of the eliminated sub-functions where necessary. Thus, these sub-functions are promoted from a per-instruction basis to a per-block basis.

Some possible modifications

[0045] It will be appreciated that many modifications may be made to the system described above without departing from the scope of the present invention. For example, the choice of sub-functions to be eliminated may be varied, according to the particular source instruction set. Equally, no optimisation may be done at all.

[0046] In another possible modification, the template initialisation process could be performed as part of the system build. In other words, part of the build process would be to compile the source templates and then to scan the resulting object code to locate the necessary information.

Claims

1. In a computer system, a method of translating source code instructions (18) into target code instructions (19) comprising the steps:

(a) analysing an existing interpreter (11) to identify sequences that implement individual source code instructions and storing those sequences as templates (14); and

(b) for each instruction in an input block of source code instructions (18), selecting an appropriate template for that source code instruction and appending this template to an output block of target code instructions(19).

2. A method according to Claim 1, including the further steps:

- (a) identifying and eliminating sub-sequences within each template (14) that implement pre-determined sub-functions;
- (b) analysing the source code block (18) to determine the net effect of the eliminated sub-sequences; and 5
- (c) planting code in the output block (19) to achieve this net effect.
3. A method according to Claim 2 wherein the sub-sequences that are eliminated include a sub-sequence for updating a program counter. 10
4. A method according to Claim 2 or 3 wherein the sub-sequences that are eliminated include a sub-sequence for clearing an overflow register. 15
5. A method according to any one of Claims 2 to 4 wherein the sub-sequences that are eliminated include a sub-sequence for performing address range checks. 20
6. A method according to any one of Claims 2 to 5 further including, for each instruction in the input block (18), determining whether the instruction requires the result of an eliminated sub-sequence and, if so, planting code in the output block (19) to supply that result. 25
7. A method according to any preceding claim further including; 30
- (a) storing fix-up information indicating constant values to be inserted into the templates (14); and 35
- (b) using the fix-up information to insert constant values derived from the source instructions (18).
8. A method according to any preceding claim wherein said step of analysing an existing interpreter comprises: 40
- (a) analysing source code (11) of said existing interpreter to identify source code template sequences (12) that implement individual source code instructions; and 45
- (b) compiling said source code template sequences (12) to generate said templates (14). 50
9. A method according to any preceding claim wherein the step of analysing an existing interpreter (11) is performed prior to run time, and the step of selecting an appropriate template (14) and appending this template to an output block of target code instructions (19) is performed at run time. 55
10. An information carrier carrying a computer program

for performing a method according to any preceding claim.

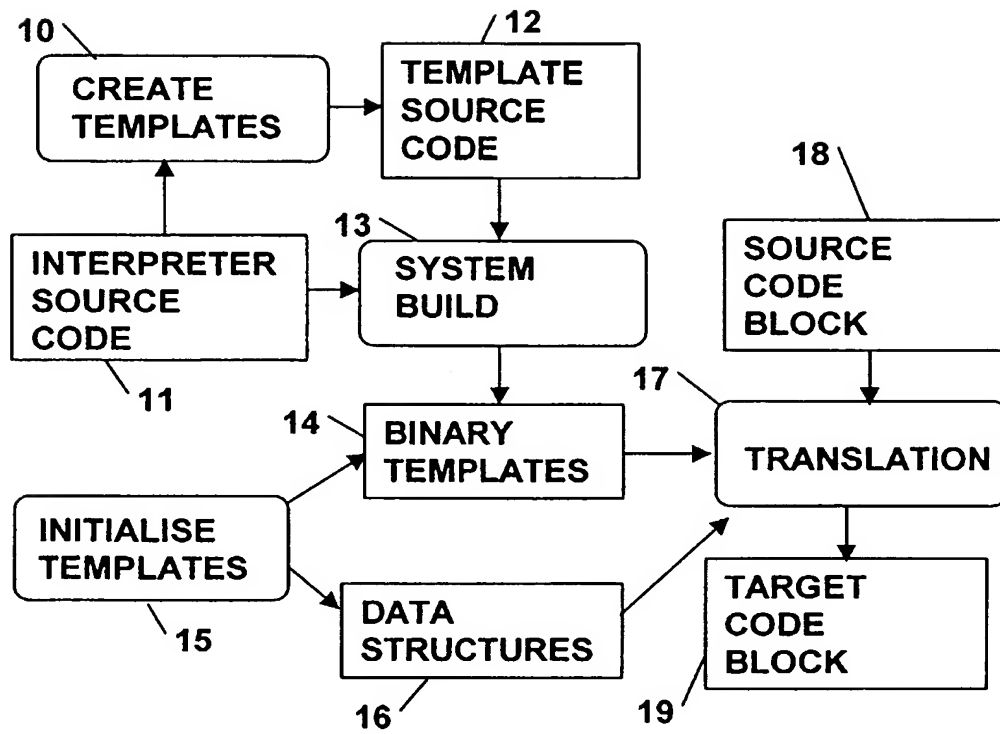


FIG.1

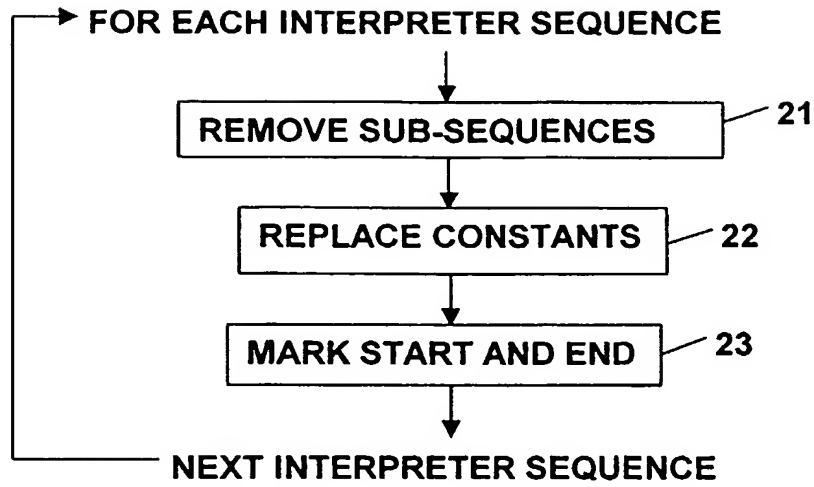


FIG.2

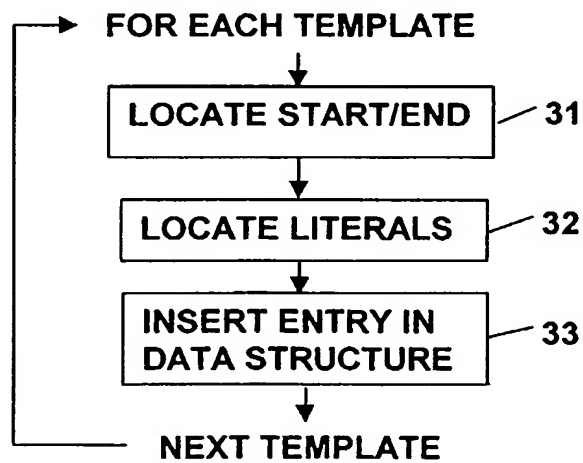


FIG.3

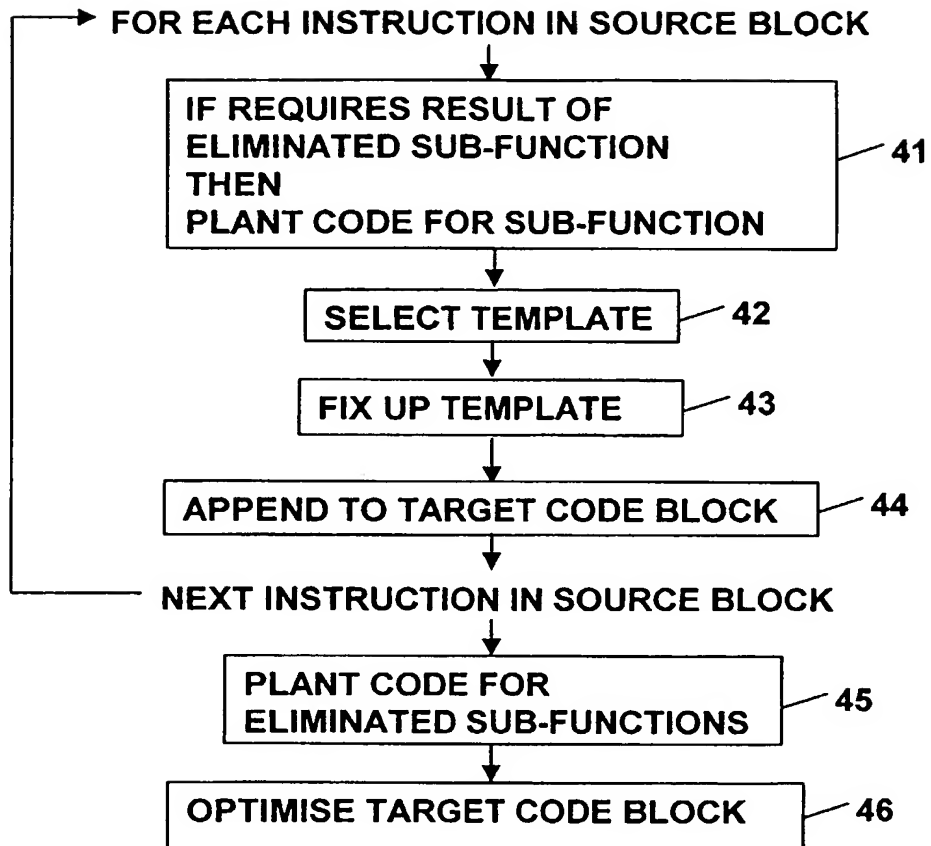


FIG.4